arXiv:1610.10064v1 [stat.ML] 31 Oct 2016

The Case for Temporal Transparency: Detecting Policy Change Events in Black-Box Decision Making Systems

Miguel Ferreira, Muhammad Bilal Zafar, Krishna P. Gummadi
MPI-SWS, Germany
{miferrei, mzafar, gummadi}@mpi-sws.org

Abstract

Bringing transparency to black-box decision making systems (DMS) has been a topic of increasing research interest in recent years. Traditional active and passive approaches to make these systems transparent are often limited by scalability and/or feasibility issues. In this paper, we propose a new notion of black-box DMS transparency, named, temporal transparency, whose goal is to detect if/when the DMS policy changes over time, and is mostly invariant to the drawbacks of traditional approaches. We map our notion of temporal transparency to time series changepoint detection methods, and develop a framework to detect policy changes in real-world DMS's. Experiments on New York Stop-question-and-frisk dataset reveal a number of publicly announced and unannounced policy changes, highlighting the utility of our framework.

1 Introduction

In modern societies, it is widely accepted that decision making systems (DMS), particularly those whose outcomes affect people's lives, need to be *transparent*. However, these decision making systems (example illustration in Figure 1) often act as black-boxes, where the precise decision making policy or function (f_{DMS}) is not known and hence the relationship between individual inputs and outputs is not clear. A number of recent studies have attempted to bring transparency to black-box decision making systems, be they driven by machines (e.g., algorithmic search and recommendation systems [9, 11]) or humans (e.g., stop and frisk decisions made by police [17, 18]). These studies attempt to reverse-engineer or infer the *decision making policy* (the function f_{DMS}) either by (i) actively auditing the system with carefully crafted inputs and analyzing the resulting outputs [9, 11] or by (ii) passively observing the inputs and outputs of the system in operation [17, 18].

The above two broad approaches to bringing transparency have their pros and cons: (i) active audits can help achieve *functional transparency*, *i.e.*, learn the behavior of the decision function for different types of inputs, but they can be expensive and might not reveal much about the system's behavior under operational conditions (where inputs are typically drawn from specific probability distributions over the input space), (ii) passive observations of the systems' inputs and outputs, on the other hand, can help achieve *operational transparency*, but they are restricted to analyzing decision function behavior only on the limited set of operational inputs seen to date.

Against this background, we make the case for a different notion of transparency that we call **temporal transparency**, where the goal is to detect *when and how the decision making policy (the function f_{DMS}) changes over time*. Note that the objectives of temporal transparency are complementary but different from those of traditional functional or operational transparency. The motivating scenarios for temporal transparency are numerous.

1. Monitoring policy change events & alerting users. Temporal transparency enables one to track

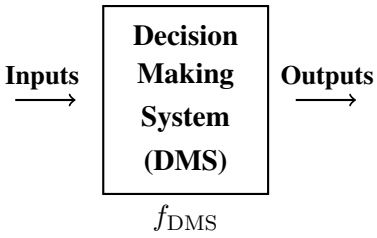

Figure 1: **The abstraction of a traditional DMS. The decision making policy (f_{DMS}) is often unknown. Efforts to bring transparency to DMS focus on inferring f_{DMS} from inputs and outputs.**

and verify when and how policies of decision making systems, such as NYPD Stop-question-and-frisk program (NYPD SQF) [1] or Facebook's newsfeed algorithm, have changed over the years [19, 20]. It would be possible to monitor whether and when an announced policy change by public or private organizations has come into effect [4, 21]. Furthermore, any unannounced (or surreptitiously deployed) policy changes can be detected and used to alert civil liberties and consumer protection groups to demand greater transparency [10]. Later in this paper, we detect several instances of announced and unannounced policy changes in NYPD SQF program.

2. Feasible when other transparency approaches aren't. Temporal transparency can be effective even in scenarios when functional or operational transparency cannot be achieved. For instance, consider the NYPD SQF program. It is not feasible to actively audit NYPD's decision making by generating artificial new inputs (*i.e.*, pedestrians in NYC). One needs to rely on passively analyzing records of stops maintained by NYPD. But, as NYPD only records data for pedestrians that have been stopped and does not record data for all pedestrians that the police are observing, it is impossibly hard to infer the decision making policy (function) in its entirety. However, as we show later in the paper, these limited records are sufficient to achieve temporal transparency, *i.e.*, robustly detect a variety of policy changes implemented by NYPD over several years.

3. Finding targets for other transparency approaches. By detecting the points in time when the decision making policy has changed, temporal transparency can help focus the more expensive traditional approaches to transparency (like active audits or passive input-output analysis) to the short period of time before and after the policy change events. Focusing transparency efforts on policy change events can help us better understand the magnitude and effects of the policy changes on the outcomes of the decision making system.

Intuitively, the basic idea behind detecting changes in decision making policy is as follows: assume we are given a time series of inputs and outputs to the system. Our task is to detect if/when the decision making policy (f_{DMS}) mapping inputs and outputs, has changed. Our intuition for detecting changes in f_{DMS} is to look for temporal changes in outputs, where inputs remain relatively stationary.

In this paper, we argue that the problem of detecting policy change events naturally fits existing frameworks for detecting changepoints in time series. Time series changepoint detection is a well-studied problem in statistics, signal processing and machine learning [5, 6, 8, 15, 22]. These studies often work with the assumption that any time series with changepoints consists of observations drawn for different statistical distributions, and at every changepoint, the distribution that the following observations will be drawn from, changes. Hence, the changepoint detection problem boils down to recovering the parameters of the underlying distributions that best explain the observations. As a by-product of this process, one also obtains a list with locations of corresponding changepoints. However, applying changepoint detection techniques on real-world datasets, subjected to noise, outliers, seasonal and weekly patterns, and different magnitudes of the detected changes, is not a straightforward task.

To tackle these challenges, we developed a framework called `Tetra` (for Temporal Transparency), that builds on Bayesian changepoint detection techniques [7, 22]. Specifically, in order to make the earlier methods robust to transitory disturbances in the observed features and aiming at detecting only significant policy shifts, we pose changepoint detection as a *maximum a posteriori* (MAP) problem

[1] https://en.wikipedia.org/wiki/Stop-and-frisk_in_New_York_City

and propose a *dynamic programming* (DP) solution. Our framework operates in an unsupervised fashion with the goal of finding the location of changepoints that best explain the underlying observations. Given an initial set of parameters to tune the sensitivity of the changepoint detection, it can return a ranked list of changepoints ordered with likelihood that a certain point indeed corresponds to a policy change. This flexibility can help the system administrator in terms of adjusting the significance level of the policy changes that are to be detected.

Applying `Tetra` on a real-world DMS, NYPD SQF program, provides interesting insights into the policy changes deployed by NYPD. Specifically, we detect several policy changes deployed by NYPD between 2006 and 2013, including changes announced publicly.

2 Detecting Policy Change Events

In this section, we outline the design of our framework `Tetra`, whose goal is to detect policy change events in a DMS.

Let $\mathbf{I}_t$ and $\mathbf{O}_t$ be the observable inputs and outputs of a DMS at time t. Let x_t be a statistic computed over $\mathbf{I}_t$ and $\mathbf{O}_t$. Consider computing x_t for a period of time $[T]$. The set of observed features collected during such time period, $x_1^T=\{x_1,...,x_T\}$, can be considered as a time series of data.

The problem at hand consist of finding the optimal set of changes—that is, the number of changes and their respective locations—which best explain the time series x_1^T. This setup can leverage time series changepoint detection frameworks. Specifically, we choose to build on Bayesian probabilistic changepoint detection setups described in [7, 22] . Adhering to the notation presented in [7], the problem above can then be formulated as:

$$\begin{array}{ll} \text{maximize} & P(\tau_1^m,m|x_1^T) \\ \text{subject to} & 1<\tau_1<\cdots<\tau_m, \\ & \tau_j-\tau_{j-1}\geq d, \\ & m\in\mathcal{M}, \end{array} \tag{1}$$

where the optimal parameters ${}^*m$ and ${}^*\tau^{*m}=\{{}^*\tau_1,...,{}^*\tau_{*m}\}$ represent the optimal number of changepoints, and their locations, respectively. $\mathcal{M}$ is defined as a symmetric set around an initial estimation of the number of changes $\hat{m}$ in x_1^T. $\hat{m}$ can be provided by the user as a part of the domain knowledge. In case the user chooses not to specify it, we consider it to be the result of computing the CUSUM chart ([6, 15]) of x_1^T and analyzing its first derivative.

Detecting significant policy regimes. Notice that, considering we are interested in detecting *significant* policy regimes, we add an additional constraint to the traditional Bayesian changepoint detection frameworks: a minimum length d of each time series segment (policy regime). The flexibility in choosing d allows for precise tunability of the framework according to a user's definition of *significant* policy regime (which may vary based on the specific application domain being considered).

Solving the MAP Problem. Given the user-defined likelihood function, $P(x_t^s|m)$, of the time series data under consideration and $P(\tau_j|\tau_{j+1})$ as the prior distribution of changepoint process, right hand side of Eq. (1) can be decomposed into:

$$\begin{array}{ll} \text{maximize} & P(m)P(x_1^T|m) \\ \text{subject to} & m\in\mathcal{M} \end{array} \left(\begin{array}{ll} \text{maximize} & P(\tau_1^m|x_1^T,m) \\ \text{subject to} & 1<\tau_1<\cdots<\tau_m \\ & \tau_j-\tau_{j-1}\geq d \end{array} \right). \tag{2}$$

By noticing the sequence of realizations (τ_1^m) form a discrete-time Markov Chain, the solution to the second term of Eq. (2) is yielded by a dynamic program, whose recurrence relation for $j\in[1{:}m]$ is dictated by:

$$\begin{array}{rcll} T(j,\tau_j) & = & \text{maximize} & P(\tau_{j+1}|\tau_j,x_1^T,m)T(j+1,\tau_{j+1}) \\ & & \text{subject to} & \tau_j-\tau_{j-1}\geq d. \end{array} \tag{3}$$

In particular, the solution to such dynamic program is given by:

$$\begin{array}{rcll} T(0) & = & \text{maximize} & P(\tau_1|x_1^T,m)T(1,\tau_1) \\ & & \text{subject to} & \tau_1\leq d. \end{array} \tag{4}$$

We fix the prior distribution of number changepoints, $P(m)$, to be a discrete Laplacian distribution (a symmetric distribution) of mean $\hat{m}$ and scale β. This choice allows the construction of the set $\mathcal{M}$ to be considered as the range of values around $\hat{m}$ which comprise a percentage α of probability mass function of $P(m)$, and the scale β of the distribution translates into the confidence in the initial estimate $\hat{m}$. The tuning parameters (α,β and other parameters regarding the analysis of the CUSUM chart) influence the shape of the set $\mathcal{M}$, and therefore the sensitivity of the setting. The joint posterior probability $P(\tau_{m-j+1}|\tau_{m-j},x_1^T,m)$ is evaluated as in [7].

Preprocessing. In order to remove underlying noise in the input time series x_1^T, and improve the reliability of the results, we apply the following preprocessing steps to x_1^T before subjecting it to changepoint detection setup outlined above:

1. **Outlier Removal.** Outliers are identified through comparison with the shifted moving average, and posteriorly removed. The size of the moving average window, as well as the threshold for outlier identification, must be adapted to each particular problem according to a user's definition of *significant* policy regime, and bearing in mind the variance in the dataset and minimum length d of each time series segment. The removal of outliers helps us ignore the extreme and noisy outputs of the DMS, providing robustness to our setup.
2. **Feature Scaling.** We scale the time series data x_1^T in a $[0{-}1]$ range, in order to simplify the setting of the tuning parameters.
3. **Filtering and Smoothing** We use a Savitzky-Golay filter [16] in order to smooth the input time series data. The parameters of the filter, its window length and the degree of its polynomial fit, are directly related with the sensitivity of the changepoint detection framework.

3 Detecting Policy Changes in NYPD SQF Program

In this section, we apply our changepoint detection framework `Tetra` on a dataset related to NYPD SQF program. The SQF program has been a subject of intense public debate since its conception [3], and went through multiple publicly announced policy changes [3, 12]. Our goal in this section is to not only check if/when the policy changes announced by NYPD were implemented but also to explore any unexpected policy changes.

To this end, we model the SQF program as a black-box DMS. We construct the time series x_t from the following observed feature: number of stops made per day under the SQF program. We assume the time series x_t to have been drawn from a Gaussian distribution. Consequently, we model the likelihood function of x_t as a Student's t-distribution, whose hyper-parameters consist of its maximum likelihood solution (MLE), computed by deploying the expectation maximization (EM) algorithm. We model the prior distribution of changepoint process to be uniform, to reflect the fact that the location of a policy change event is independent of the location of the previous one. Finally we specify the minimum length d of each time series segment to be 15.

We deployed `Tetra` on the stops made during the years 2006 to 2013 (inclusive). The complete records of the stops made under SQF program are made publicly available at the official website of NYC. [2] Our framework detected a total of 31 changepoints. Since the number of detected changepoints is considerably large, we systematically analyzed each of the changepoints. As a result, we were able to separate the changepoints into following categories (listed individually for each year in Table 1):

1. Seasonal patterns. These changepoints correspond to slight drops in number of stops made each day around mid-year (summer) and close to the end of the year (winter). This pattern persists for almost all of the years considered for the analysis. 16 out of the 31 detected changepoints fall under this category.

2. Unusual input changes. These changepoints *potentially* correspond to unusual changes in everyday pedestrian population of NYC. For example, we detect a changepoint on October 29, 2012, marking a consistent drop in number of stops made per day un-

[2] nyc.gov/html/nypd/html/analysis_and_planning /stop_question_and_frisk_report.shtml

Year	Seasonal		Unusual inputs	Policy	
	S	W		A	UA
2006	1	1	–	–	–
2007	2	1	1	–	–
2008	1	1	1	–	–
2009	–	2	–	–	1
2010	1	1	–	–	2
2011	1	1	2	–	1
2012	2	–	2	1	1
2013	–	1	–	–	3

Table 1: **List of detected changepoints from January 01 to 2006, to December 31, 2013. S—Summer; W—Winter; A—Announced, UA—Un-announced.**

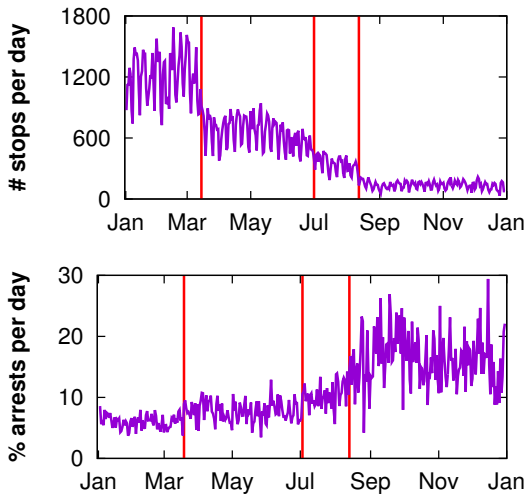

Figure 2: **Changepoints detected in NYPD SQF data from January 01, 2013 to December 31, 2013.**

til the next changepoint on November 10, 2012. This drop is most probably due to Hurricane Sandy and its aftermath [1]. In fact, on the day when the changepoint occurred—October 29, 2012—the number of stops made over the city is merely 193, as compared to an average of 1147 stops per day for the previous week. Similarly, a changepoint marking an increase in number of stops per day on September 22, 2011 could potentially correspond to Occupy Wall Street Movement, that started on September 17 [2]. In total, 6 out of 31 changepoints map to this category.

3. (Un)Announced Policy changes. The changepoints that correspond to neither of the above two categories were likely caused by policy changes implemented by NYPD (because they cannot be explained by input changes). For example, we detect a drop in the number of stops made per day starting March 26, 2012. This change is in fact a consequence of a *publicly announced* policy change implemented by NYPD, where, 'increased training' and staffing in 'high impact' zones results in an overall decline in number of stops [12]. Detection of this changepoint highlights the utility of our framework in verifying the policy changes announced by the governing entities.

Next, we focus on analyzing changepoints that do not map to a publicly announced policy change. In particular, we focus the year 2013. The changepoint detection framework yields 3 un-announced changes for this year. Figure 2 (top panel) shows the number of stops made per day and the detected changepoints (in the form of vertical lines). Remarkably, this series of changepoints correspond to three *abrupt* policy changes which successively brought down the number of stops per day to eventually 10% of the stop rate at the beginning of the year. It is important to note that the 2013 SQF program was subject of intense debate during the 2013 Mayoral Election campaign, with a major candidate denouncing it [14] and a court stating that the SQF policy violated the constitutional rights of the citizens [13]. Consequently, these variations are likely to be associated with *un-announced* policy adjustments resulting from these events.

In addition to studying the number of stops, we also analyzed the percentage of stops leading to arrests per day in 2013. The changepoint analysis framework detects three changepoints presented in Figure 2 (bottom panel), close to the changepoints detected in the stop-rate analysis. This clear mapping between the changepoints yielded by both observed features reveals a systematic change in SQF policy by NYPD, indicating that the policy change did not just concern the number of stops per day, but also the *nature* of the stops.

4 Conclusion and Future Work

In this paper we made the case for temporal transparency where the goal is to detect when and how the DMS policy changes over time. We built a framework `Tetra` using prior advances in Bayesian changepoint detection. Applying `Tetra` on a real-world dataset shows that it can systematically detect possible policy change events in practice. In the future we hope to generalize our framework to apply it to a broader range of real world DMS's. More specifically, we plan to address the following points:

1. The current implementation relies on an 'offline' setting that needs access to the whole time series data to detect possible changepoints. Hence, the framework cannot be deployed on streaming datasets, where one might want to detect changepoints on the fly, *e.g.*, Facebook newsfeed algorithm. We are currently expanding it to incorporate an 'online' setting in order to cater to such scenarios.

2. As shown in Section 3, analyzing the structure of policy changes by jointly considering multiple observed features (number of stops, percentage of stops leading to arrests, in parallel) can provide more insights into how the DMS interplays with different features, hence revealing more information about policy changes. To address this point, we plan to generalize our framework to multi-variate feature spaces.

References

[1] `http://www.marketwatch.com/story/cuomo-orders-nyc-transit-system-to-shut-down-2012-10-28`.

[2] `http://occupywallst.org/about/`.

[3] Stop-and-frisk in New York City. https://en.wikipedia.org/wiki/Stop-and-frisk_in_New_York˙City.

[4] Al Jazeera America. Monitor: Changing NYPD stop-frisk practices a challenge. `http://america.aljazeera.com/articles/2016/2/17/monitor-changing-nypd-stop-frisk-practices-a-challenge.html`, February 2016.

[5] D. Barry and J. A. Hartigan. Product partition models for change point problems. *The Annals of Statistics*, pages 260–279, 1992.

[6] M. Basseville, I. V. Nikiforov, et al. *Detection of abrupt changes: theory and application*, volume 104. Prentice Hall Englewood Cliffs, 1993.

[7] P. Fearnhead. Exact and efficient bayesian inference for multiple changepoint problems. *Statistics and computing*, 16(2):203–213, 2006.

[8] P. Fearnhead and Z. Liu. On-line inference for multiple changepoint problems. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 69(4):589–605, 2007.

[9] A. Hannak, G. Soeller, D. Lazer, A. Mislove, and C. Wilson. Measuring price discrimination and steering on e-commerce web sites. In *Proc. IMC'14*.

[10] Huffington Post. Facebook Just Made A Pretty Awkward Change To Your Profile. `http://www.huffingtonpost.com/entry/facebook-intro-work_us_5769831e4b015db1bca97c9`.

[11] M. Lécuyer, G. Ducoffe, F. Lan, A. Papancea, T. Petsios, R. Spahn, A. Chaintreau, and R. Geambasu. Xray: Enhancing the web's transparency with differential correlation. In *USENIX Security'14*.

[12] New York Post. Major decline in NYPD stop-frisks. `http://nypost.com/2013/02/09/major-decline-in-nypd-stop-frisks`.

[13] New York Times. Judge Rejects New Yorks Stop-and-Frisk Policy. `http://www.nytimes.com/2013/08/13/nyregion/stop-and-frisk-`

practice-violated-rights-judge-rules.html.

[14] Newsweek. Did Bill De Blasio Keep His Promise To Reform Stop-and-Frisk? http://europe.newsweek.com/did-bill-de-blasio-keep-his-promise-reform-stop-and-frisk-266310.

[15] E. Page. Continuous inspection schemes. *Biometrika*, 41(1/2):100–115, 1954.

[16] A. Savitzky and M. J. Golay. Smoothing and differentiation of data by simplified least squares procedures. *Analytical chemistry*, 36(8):1627–1639, 1964.

[17] R. S. Sharad Goel, Justin M. Rao. Precinct or Prejudice? Understanding Racial Disparities in New York City's Stop-and-Frisk Policy. *Annals of Applied Statistics*, 2015.

[18] C. Simoiu, S. Corbett-Davies, and S. Goel. Testing for Racial Discrimination in Police Searches of Motor Vehicles. *SSRN abs.2811449*, 2016.

[19] The New York Times. Facebook to Change News Feed to Focus on Friends and Family. http://www.nytimes.com/2016/06/30/technology/facebook-to-change-news-feed-to-focus-on-friends-and-family.html, June 2016.

[20] Time. Here's What Facebook's Big New Change Really Means. http://time.com/4387908/facebook-change-news-feed-update/, June 2016.

[21] C. S. Times. Chicago police and aclu agree to major changes in stop-and-frisk policy. http://chicago.suntimes.com/politics/chicago-police-and-aclu-agree-to-major-changes-in-stop-and-frisk-policy/, August 2015.

[22] X. Xiang and K. Murphy. Modeling changing dependency structure in multivariate time series. In *Proc. ICML'07*.